

### Appendix 3: Zero-Casualty War Scenario

A massive conflict has broken out between countries A and B in Asia. Country B has invaded city “a,” whose population is historically 40 percent of Country B origin. The attack includes possible use of chemical and biological agents. Country B has occupied the border valley leading up to city “a.” The inhabitants of city “a” are locked in, but refugees from the surrounding region are pouring out into neutral Country C. Reports of atrocities and suffering within and around city “a” abound.

The U.S. has sent in humanitarian aid. The United Nations has imposed economic sanctions, which have not stopped the attacks, and has subsequently issued an ultimatum to Country B to withdraw its forces. The UN has now approved the use of force.

In preparation for conflict, the U.S. has deployed networks of sensors in and about the conflict region. These include chemical sensors (vehicle exhaust fumes, urine, chemical agents, etc.), broad-spectrum acoustic sensors, seismic sensors, video sensors, and imaging sensors. Some are mobile. The U.S. has coordinated with other regional powers for cooperative actions, and is preparing to engage in conflict promising — to the U.S. Congress — minimal casualty war.

The U.S. tasks the Army’s FFCS (Future-Future Combat Systems) to:

- ◆ Contain or stop attacks on the civilian population and the refugees in the border regions
- ◆ Remove Country B’s forces from Country A with minimal civilian casualty
- ◆ Establish a Zone of Separation between Country B and Country A
- ◆ Ensure that Country B’s atrocities and aggression are displayed across the Internet and the media
- ◆ Conduct hand-over to follow-on forces when so ordered

This mission is to be achieved through a set of precision actions whose duration should not exceed 5 to 10 days.

To achieve these goals, the U.S. performs the following actions (this is an incomplete list):

- ◆ Establishes a forward operating base in Country C with reach-back to the continental U.S. (CONUS) over satellite (or whatever other communication links are reliably available)
- ◆ Deploys Unpiloted Aeronautical Vehicle (UAV) Networks (swarms of coordinated and cooperating UAVs ) controlled via satellite and other relay links, and collaborating with each other via cross-link
- ◆ Establishes identification and tracking of each military air and ground vehicle in the area of operation
- ◆ Dispatches combined Air Cavalry and FFCS units from U.S. bases in Hawaii

The sensors are deployed via cruise missiles from ships in the Indian Ocean. The forward operating base is dispatched from Europe via transport planes, as are the UAV swarms. The UAVs consist of a mix of fixed- and rotor-wing aircraft, with altitudes of operation varying from tree-top level to 65,000 feet.

The forward operating base begins to perform tactical reconnaissance. Terrain, street, and building information are updated based on visual and acoustic information from the UAVs. Signals Intelligence data are gathered from the UAVs and relayed to the forward operating base for analysis and correlation. The acoustic, seismic, and visual signatures of each of the hundreds to thousands of motorized vehicles are cataloged with each vehicle being identified as military or civilian, and military vehicles are identified as hostile or friendly. Air Defense Artillery and Surface-to-Air missile sights are identified. Automated analysis of visual information provides data about the approximate number and locations of dismounted troops, enemy command posts, and command vehicles. Weapons and supplies stores are cataloged as identified. Visual, chemical, and acoustical indications of weapons fire are all enunciated within the forward operating base and video of that region is either initiated or tagged. Personnel within the forward operating base monitor tagged video for “newsworthy” clips and forward as appropriate to news media. Forward operating base data are relayed to CONUS. Video streams and live sensor reports may be relayed to CONUS at the initiation of CONUS or forward operating base personnel.

While en-route aboard transport planes, FFCS and Air Cavalry units begin to monitor vehicle-tracking information and dismounted troop movements. Although FFCS vehicles are disbursed among different transport vehicles, inter-vehicle communication is accomplished via air-to-air cross-links and each transport vehicle has high-rate connectivity to the satellite resources. Up-to-date terrain, street, building, and weather information is loaded into FFCS and Air Cav onboard databases via satellite from CONUS and the forward operating base. As the tactical picture emerges, officers in the forward operating base, CONUS, and aboard the transport planes collaborate to finalize operational plans. One or more FFCS units (“Cells”) are assigned to each of the task forces that are to execute the following assignments:

- ◆ Task Force Alpha: Prevent the Country B troops occupying the border valley from moving further into Country A
- ◆ Task Force Bravo: Cut off the logistics supply chain to Country B's forces
- ◆ Task Forces Charlie and Delta: Flank the Country B troops in the valley
- ◆ Task Forces Echo, Foxtrot, Golf, and Hotel: Encircle city “a” and use progressively lethal robotic means to disable or destroy Country B assets within the city. Coordinate with Special Operations Forces tasked to covertly evacuate Country A leaders from city “a” during the initial hours of the FFCS operation.
- ◆ Task Force India: Establish and occupy battle position to provide Air Cavalry support to other task forces
- ◆ Task Force Juliet: Reinforce forward operating base and establish refugee and enemy prisoner of war facilities

While still en route, task forces develop unit plans and assume responsibility for specific roads, intersections, and vehicles. Using intervisibility calculations derived from the terrain

data, FFCS units determine where to place robotic direct fire vehicles for optimal range fans and where to place robotic indirect fire vehicles, sensor vehicles, infantry carriers, and command and control vehicles.

Upon arrival at Country A, task forces deploy under cover of darkness, with some task forces being parachuted into position while others are dispatched from the forward operating base. Task forces deploy their organic UAV support, and the command and control vehicles establish processed sensor data and live video support as required from the UAVs under the control of the forward operating base.

Three hours before dawn, the Special Operations Forces enter the city and rendezvous with Country A's leaders who are to be extracted.

One hour before dawn, task forces launch indirect fire missiles to destroy all known Country B supply and weapons dumps in Country A. The Special Operations Forces depart the city during the ensuing confusion. At the same time, additional indirect fire missiles are targeted at Country B command posts and vehicles closest to the city. These missiles are instructed to loiter above their targets, above the range of enemy fire. The UN ultimatum is repeated to Country B and the surrender of Country B forces is solicited.

In response to the ultimatum, Country B forces within the city begin to fire on civilians. Muzzle-flash sensors identify the source of the firing, and loitering missiles are tasked to eliminate these vehicles. The call for surrender is repeated and Country B requests to withdraw. The task forces oversee the retreat of Country B troops and establishes the Zone of Separation along the border between Country A and Country B.

This mission succeeds with minimal casualties as a result of:

- ◆ Large scale sensor fusion
- ◆ Large scale target identification and tracking
- ◆ Large scale video acquisition, transmission, analysis, and routing
- ◆ Robotic command and control of surface and aerial forces
- ◆ Rapid insertion of overwhelming force

A key component of this mission is the transmission of critical, sensitive information over reliable, secure networks. The networks must support command and control as well as tactical operations. Functional networking requirements include the:

- ◆ Ability to disseminate information appropriately over high to low bandwidth links, assigning appropriate flows to appropriate links when multiple paths exist simultaneously
- ◆ Accommodation of voice, data, video, audio, still images, etc.
- ◆ Ability to provide high assurance, real-time closed-loop distributed control in a tactical environment while on the move in harsh terrain
- ◆ Seamless integration of heterogeneous component networks, including FFCS nets, other sensor nets, and command structure communications

- ◆ Provision of hierarchical views of the network structure, depending on the "user": forward operating base, CONUS, Task Force commanders, and FFCS Cell team leaders all have access to common views of the tactical situation, but typically operate at different levels of the hierarchy to enhance their effectiveness

#### **Appendix 4: Deeply Networked World/SWARMS Scenario (Smart World Airforce Repair and Maintenance System)**

We imagine in this future scenario the truly modernized Airforce, highly efficiently managing its plane maintenance and parts scheduling. The Airforce has recognized that they have several problems. First, in order to save money, the Airforce has learned from industry that they want to minimize their stock of parts and provide just-in-time delivery of repair parts. Second, the places at which the Airforce will need parts vary because the planes travel around the world and it is more efficient to have the parts at the locations where the planes will need them rather than requiring that the planes return to central repair facilities. The reasoning behind this is that the planes may be needed in action, and moving them far from their military arena for repairs or worse for routine maintenance is undesirable.

In that future time, we find a world where networking and networked devices have become more broadly and deeply deployed. In this world it is difficult to find places that are off “the net” because locations have available some form of wired, wireless, or other networking technology. Furthermore, the network is able to provide service to an extremely large number of devices essentially co-located. As a society, we have bought into a truly pervasive computing environment.

It has been recognized that humans cannot and should not be responsible for the detailed management of such a deeply networked world, nor for managing the quantity and heterogeneity of devices. The network must be not only self-organizing, but self-healing. To do this, there must be capabilities for measuring behavior, evaluating that behavior, and then either masking or correcting it, when problems arise. An “agent” model must be deployed, in which agents not only function correctly as individuals, but also understand how to collaborate with each other effectively in order to make higher level decisions than any individual agent might make. Both individuals and groups providing higher level, composite functions must be responsive to societal policy constraints that may change or evolve with time. The developers of this technology call it the *smart world*.

The truly modernized Airforce has bought into the *smart world* model and has determined that every repair component and parts depot will be “smart.” A part will know where it is. A depot will know how many of each kind of part it has and have a model of what parts are needed based on reports about the schedule of arriving planes. The Airforce has gone further than that. Every part, when installed in a plane, is also introspective. Each knows how well it is functioning, and can predict when it will need to be replaced. Beyond that the composite systems not only integrate over all their parts, but also have a higher level understanding of the emergent system. Agents will query, collate, and manage the system. The architecture is called the Smart World Airforce Repair and Maintenance System or SWARMS. SWARMS can predict when and where specific repairs will be needed. At the level of the whole plane, it understands flight schedules, using this information to plan where and when work should be done. SWARMS will inform the global inventory system, which ensures that by the time a plane arrives at a destination the appropriate parts are there, with enough information that the repair can be made.

The Air Force is one of many branches of the government that have bought into the new “smart world” model. As a result, although many *smart worlds* are evolving, they will share some common network and services infrastructure, such as backbones. Furthermore, in some situations resources need to be shared among *smart worlds* (for example, requisitioning parts from outside the Airforce and interoperation under circumstances such as civilian crises in which the military helps provide support). These often reflect situations where coordination is needed or it is necessary to collaborate on the use of higher level resources.

In a military operation flexibility through over-provisioning is needed because predictive and planning capabilities will not be perfect and equipment reliability is critical.

Two objectives of the opposing force in this situation are espionage and sabotage. They would like to know what sorts of planes will be deployed at what locations and when. In addition, if possible they would like to cause those planes to be disabled.

The opposing force believes that SWARMS may be able to help them. Many of the individual parts and perhaps even collections of them, as well as the agents, were designed specifically to report on their state and make predictions about themselves, indicating when and under what conditions repairs or replacements will be needed. The agent-based network infrastructure may also help. The network is populated by network management agents that are collecting information, making decisions about network resources and behavior, and reporting as needed. Since these agents are “in the core of the network” they have to be trusted. However, different communities and individuals may have quite contradictory constraints on their operation and behavior, as well as differing policy controls. The opposing force believes that the combination of the smart network management world and SWARMS will allow for both the espionage and sabotage activities giving them the upper hand in the engagement.

In the context of SWARMS and this scenario there is a complex set of issues such as trust, assurance, and security, including privacy, authenticity, authorization, and denial of service. For example, one can consider that the “Unified Smart World” will consist not only of the Airforce’s planes, but large numbers of additional planes and systems from the military and civilian communities. The motivation for ensuring smart behavior may range from economic forces to military security to life and death issues. If one of the Airforce’s planes does not make it to a destination because of a part failure, this may cause not only loss of the plane, but also loss of life. To focus in further, one can ask what needs to be in the network and what can be pushed to the edges, recognizing that the true edges may be “dumb” devices with limited functionality and adaptability. Finally, one can consider the problem of providing assurance and trust in emergent systems.

Scenario networking elements:

- ◆ Deeply networked systems
- ◆ Security including authentication, authorization, and privacy
- ◆ Smart network management
- ◆ Network performance measurement

## **Appendix 5: Crisis Management Scenario**

In 2010, NOAA began a five-year deployment of its Ultra Doppler/SAR radars across “Tornado Alley” in the central U.S. In the following year, DoD received “dual use” approval to supply early fire warning bulletins from its “staring” missile launch detection satellite systems. And in 2012, NASA orbited Firesat, capable of providing twice-a-day high resolution multi-spectral, multi-instrument views of forest fire activity around the world.

The year is now 2015, and a “perfect” fire and tornado season has descended on the U.S. Hot, dry Santa Ana winds have come to the West Coast of the U.S. with a vengeance, from San Diego all the way up to the Pacific Northwest. In the Central States, destructive twisters are beginning to form along nearly every low-pressure system that sets up across Oklahoma and Texas.

On the day of crisis, DoD reports dozens of new fires from a single highly charged lightning storm. By mid-day, the early morning Firesat images and data have been processed and disseminated to hundreds of state and Federal agencies. Hot-spot data are combined with vegetation cover and dryness sounder models to produce detailed next-24-hour maps for the worst hit regions. These maps and digital models are disseminated instantly to government command and control centers and are spotcast to individual homes and businesses most in danger.

Department of Interior and other Federal agency supercomputers begin “nowcasting” the present and predicted tracks of the worst fires. Their sophisticated models take into account detailed wind and temperature forecasts transferred in near real time from NOAA’s National Weather Service computers as well as vegetation burn models built from satellite data and from digital elevation model data in U.S. Geodetic Survey computers. These models and data are transferred directly to Federal and state fire management command and control centers.

The models and nowcasts are also transferred by satellite communications to the forest fire field units, which return validation and update information. This field information, along with real-time atmospheric, chemical, and other environmental data from sensors deployed throughout the area – both in-situ microsensor platforms deployed in advance as well as self-contained, self-powered sensors dropped from aircraft that same morning – are continuously integrated into the nowcast models. Customized warning and evacuation messages are automatically provided to all the homes and businesses in the areas.

On the same day, tornadoes are touching down across a broad path in the central U.S. Data from the Ultra radars are continuously fed to NOAA supercomputers, which nowcast the locations and severity of the funnel clouds and touchdowns with half-mile accuracy. Within two minutes, every home that has a 60 percent or better chance of being in the path of a tornado receives a computer-generated call on its telephones and cell phones to take cover. Houses that have been automated go into protective mode, shutting off natural gas feeds, closing drapes, and taking other actions that they have been designed to do. Non-automated houses are shut down remotely by their owners from work, using secure Web browsers to access sensors and actuators plugged into their home networks. Police and fire units as well

as all individual homes in each affected area receive detailed nowcast maps of the tornado tracks and touchdowns in their neighborhoods. This allows the authorities to mobilize resources safely to provide emergency medical support to each affected site within two minutes after touchdown.

Immediately after each tornado hits, with power lines down and communication infrastructures demolished, emergency mobilization forces are directed by computer into the target zero area. The emergency units are set up quickly, establishing a high-performance field network instantly capable of local area and remote communications, using truck-based wireless technologies tied into regional networks via high performance satellite communications. Mobile whole-body scanners, sophisticated medical instruments, and mini chemical analysis labs are plugged into the network. This allows a formidable collection of medical specialists, data and information resources, and analysis facilities to be called on instantly to support the onsite paramedics. Command and control units have instant high-performance network access to all needed statewide and Federal resources. Within five minutes after touchdown, the President appears on television talking via two-way videophone to the area commander and a survivor at the worst touchdown site, declaring the region a disaster area and promising that the full resources of the Federal Emergency Management Agency will be made available to help the victims and promote a speedy recovery.

Scenario networking elements:

- ◆ Input from heterogeneous sensors
- ◆ Large-scale on-line modeling, QoS, and bandwidth requirements
- ◆ Automated delivery to diverse sites
- ◆ Self-contained small sensors
- ◆ Self-organization of networks
- ◆ On-line medical consultation capabilities



## Appendix 6: Collaboration Scenario

Collaborative problem solving and decision making is a fundamental aspect of Dr. Clotho's research, which includes a large number of national and international partners and collaborators. Furthermore, much of the research is focused on domains in which complex tasks must be performed in environments that are largely inaccessible to human beings.

Dr. Clotho indicated that, "Collaboration can be much more fruitful if we can conduct, from the comfort of our labs, scientific experiments where all resources are located remotely, and yet all phases of the experiments are so well orchestrated that it appears a local endeavor to all participants. This should be possible, not only for immediate collaborators but also for everyone else that cares to join and participate in the research."

Dr. Clotho stated, "Current systems do not allow true collaboration and thwart natural, intuitive styles of interaction. They are passive in design and are difficult to adapt to particular work patterns. The services they provide are primarily fragmented and non-interoperable, which quickly makes them obsolete. How about security and the need to incorporate new applications in the collaboration space quickly and effectively?"

"The scientific environment will include a large number of sensors and robots with varying capabilities, small enough to be embedded into the natural environment with minimum disturbances. The simple ones may be able to sense, compute, and act. The more sophisticated ones may be able to carry out more complex tasks related to detailed physical monitoring and manipulation. However, only through deployment of dense spatial sensing and the coordinated effort of a large number of these nodes can I explore my physical world and carry out my research goals.

"These low-power nodes, with limited communication bandwidth, need to understand local conditions and together collaborate to identify and monitor global environment conditions. They need to be able to form collaborative teams to accomplish complex tasks, such as surveying sites of scientific interests and coordinating to overcome potential problems and failures as they occur. The right level of coordination must be selected dynamically depending on the task at hand and the feedback from the scientists. For tasks where limited coordination is sufficient, the nodes can form teams and negotiate roles in a local yet distributed fashion. For tasks in which tighter coordination is required, the nodes can take a more global approach, potentially seeking the help of the scientists, in assigning roles."

"The problem," continued Dr. Clotho, "is that no single node has global knowledge and their capabilities preclude any centralized coordinated global sharing of state. So do not talk to me about centralized knowledge or control. It is simply not achievable for the large number of devices that I will be deploying at possibly fine granularities. It will be much easier for my colleagues and myself if coordination schemes that are inherently distributed and based on localized inputs, algorithms, and outputs were made available to us. Take the case of data collection, for example. This task is almost insurmountable in my small-scale simulated testbed. I cannot even begin to imagine the inevitable implosion of data that I will

be facing from the need to continuously monitor, at high resolution, the physical world I am interested in exploring. My task would be much easier if it were possible to execute local correlation and possible aggregation of data inside the network before I collect and process the data at the desired level of granularity.”

Then in reference to her collaboration with her direct national and international peers, Dr. Clotho added, “Facilitating the interaction and collaboration among the large number of limited devices addresses only one aspect of my problem. We also need collaborative support at the scientist's level. We need tools for pro-active and dynamic support so that the system can behave as a problem-solving partner, adapting to that scientist's work pattern, potentially providing advice to advance collaboration. Each scientist should be able to augment the reality of the physical world with virtual worlds to allow the automated inclusion of ‘what if’ scenarios and conduct studies tailored to specific interests.”

These requirements pose unprecedented challenges for future scalable, robust distributed system design. Add to that the unpredictability brought about by the micromobility of the devices and you will find that requirements and constraints cannot be addressed by current networking and distributed coordination technologies. Not only are traditional protocols too heavy-weight, but the building blocks of our current distributed systems and networking, namely layering, abstraction, and modularity, become questionable. The framework within which future collaborative research will be carried out must be flexible enough to deal with unpredictable and emergent changes, meet hard real-time constraints, and handle asynchronous events as they occur.

## Appendix 7: Networked Medical Care Scenario

After eating a heavy Tex-Mex dinner, a middle-aged man is at home watching television one evening when he begins to suffer chest pains. Is he having a heart attack or is it just heartburn? He goes to his medicine cabinet where he finds a rod-like device that he presses against his bare chest. A yellow light comes on. The light soon turns green. The man presses the big red button. A few minutes later the telephone rings. A voice informs the man that he is having a heart problem and that an ambulance has been dispatched to his home.

The ambulance arrives to take the man to the hospital. On the way an EMT attaches the man to the vital signs and cardiac monitoring equipment. A cardiologist, who is at home but on duty, can be seen on the monitor in the ambulance. She has done a network search and has access to the patient's complete multi-media medical record. The cardiologist can see the patient and guides the EMT through a cardiac evaluation that includes heart sounds, a complete EKG and an echocardiogram. Based on her evaluation of the data received from the evaluation and her view of the patient, the cardiologist orders an angiogram to be performed as soon as the patient arrives at the hospital. The hospital angiogram team is assembled, scrubbed, and ready when the patient arrives.

The patient is prepared for the angiogram as soon as he arrives at the hospital. The procedure is performed. The hospital angiographer thinks that he has detected a cardiac anomaly but is unsure so he calls in a consultant. The consultant is reached at home and agrees to look at the angiogram on his home computer. When the consultant accesses the angiographic data over the network, the hospital angiographer receives a warning from the network that the video screen being used by the consultant does not meet the standard deemed necessary for angiogram interpretation. The consultant does not see the anomaly so, despite the warning from the network, it is decided that bypass surgery is necessary. Preparations are made to perform the surgery in the morning.

Bypass surgery begins. It appears to be a routine case until the surgeon realizes that the patient's cardiac vasculature follows a rare anatomical anomaly for which the surgeon has no experience. From the operating room the surgeon is able to search a 3-D anatomical library for a similar case. He contacts the colleague responsible for the case he has found. The consultant is able to see the patient's cardiac vasculature and advise on the proper way to successfully complete the operation. He might even occasionally take control of the haptic enabled surgical robot in order to help the resident surgeon through some of the more difficult or unfamiliar procedures.

Scenario networking implications:

### *Security*

First and foremost, the network must be secure and must meet legal standards for medical data privacy and security, i.e., currently the proposed Health Insurance Portability and Accountability Act rules. The network must be able to securely carry stored data in a wireless environment.

### *Scalability and High Assurance Networking*

To support the cardiologist at home, the wireless channel must provide real-time video and real-time data signal display. Although this application may tolerate a fair amount of latency, it will not tolerate jitter and the video, audio, and data channels must be synchronized.

### *End-to-End Performance, Intelligent Networking*

The angiogram procedure identifies the need for end-to-end knowledge of the network data path including the display devices at the ends. An intelligent, scalable network should automatically find the work-around for the quality of service problem, but it must report the problem if a work-around is not possible.

### *Assured Real-time Service*

For the bypass surgery scenario, the surgeon needs to retrieve 3-D image data sets, each of which may be several gigabytes in size. The consultant must be able to view the live procedure in real time and accurately guide the surgical robot through its haptic controls. This requires a network that is able to operate at high bandwidth with minimal latency and minimal jitter while maintaining the security and integrity of the transmissions and the privacy of the patient data.

## Appendix 8: High Energy Physics Scenario

Never before has the scientific mission of particle physics research been so dependent on state-of-the-art information technology. Collaborations of hundreds to thousands of physicists and engineers are formed to create accelerators, detectors, and analysis systems with a productive life of tens of years. These analysis systems form a complex and widely distributed “fabric” of computing and storage resources.

The non-deterministic nature of quantum physics, uneasily understood during the last century, inevitably requires the measurement and analysis of billions of particle interactions to observe and understand fundamental processes. Particle physics experiments have pushed against the limits of technology, electronics, computing, and networking for decades. Detectors with millions of channels, each recording precise amplitudes with a resolution of picoseconds, have in the course of 40 years succeeded detectors with a few single-bit “yes/no” measuring devices. Information flows from such a detector at up to a terabit per second and must be drastically filtered in real time because of limited storage, analysis, and networking facilities.

The Large Hadron Collider (LHC) experiments at the European Organization for Nuclear Research (CERN) will rapidly reach tens of petabytes of stored data under intense analysis. The design, construction, and data analysis for an experiment require the combined intellect and dedicated work of international collaborations. However, technological limitations on the storage, transmission, and analysis of data impose difficult, even dangerous choices. For example, the LHC experiments expect to be able to record and share over networks less than one millionth of the collisions they observe. This draconian real-time selection will necessarily have to be optimized for “somewhat expected” new discoveries rather than the “totally unexpected” ones that are the dream of every scientist.

Even after the draconian selection, LHC collaborations will face the challenge of empowering thousands of geographically distributed physicists to use their intellect and wisdom to derive physics insight from tens of petabytes of data. Although the raw cost of bandwidth is no longer a crippling impediment, the end-to-end performance of applications is often unacceptable. Success will rely on middleware research to support data-intensive, worldwide collaborative science that is only beginning. A minimum requirement is the location-independent ability to analyze data to empower all of an experiment’s physicists to work collaboratively on databases, growing to tens of petabytes in 2005-2010, using all computing resources to which they have access.

However a qualitative change in the way research is performed would be enabled if we could free the real-time selection of data from the constraint of a single filter system with selection algorithms decided by committees. The availability of networks with end-to-end terabit performance could make this possible, but speed alone is not enough.

The data acquisition and filtering systems might profitably become geographically distributed and operate as highly parallel, largely asynchronous data flow systems. The terabit systems that will become operational in 2005 for the LHC will include a multi-terabit

capacity switching fabric, but individual data acquisition nodes and filtering nodes will communicate at gigabit speeds. In addition to the substantial bandwidth requirement, challenges include:

- ◆ The multicast service required when more than one remote filtering center is available
- ◆ Achieving adequate error rate and robustness without *ever* allowing the implementation of the “wild idea” to impact the detector-site data acquisition system